

Recursively Conditional Gaussian for Ordinal Unsupervised Domain Adaptation

Xiaofeng Liu^{1†*}, Site Li^{2†}, Yubin Ge³, Pengyi Ye¹, Jane You⁴, Jun Lu⁵

¹Harvard University, Cambridge, MA, USA.

²Carnegie Mellon University, Pittsburgh, PA, USA.

³Dept. of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL, USA.

⁴Dept. of Computing, The Hong Kong Polytechnic University, Hong Kong.

⁵Beth Israel Deaconess Medical Center, Harvard Medical School, Boston, MA, USA.

[†]Contribute Equally, ^{*}Corresponding Author: liuxiaofengcmu@gmail.com

Abstract

The unsupervised domain adaptation (UDA) has been widely adopted to alleviate the data scalability issue, while the existing works usually focus on classifying independently discrete labels. However, in many tasks (e.g., medical diagnosis), the labels are discrete and successively distributed. The UDA for ordinal classification requires inducing non-trivial ordinal distribution prior to the latent space. Target for this, the partially ordered set (poset) is defined for constraining the latent vector. Instead of the typically *i.i.d.* Gaussian latent prior, in this work, a recursively conditional Gaussian (RCG) set is adapted for ordered constraint modeling, which admits a tractable joint distribution prior. Furthermore, we are able to control the density of content vector that violates the poset constraints by a simple “three-sigma rule”. We explicitly disentangle the cross-domain images into a shared ordinal prior induced ordinal content space and two separate source/target ordinal-unrelated spaces, and the self-training is worked on the shared space exclusively for ordinal-aware domain alignment. Extensive experiments on UDA medical diagnoses and facial age estimation demonstrate its effectiveness.

1. Introduction

Deep learning is typically data-starved and relies on the *i.i.d.* assumption of training and testing sets [5, 32]. However, the real-world deployment of target tasks is usually significantly diverse, and massive labeling of the target domain data can be expensive or even prohibitive [29]. To address this, the unsupervised domain adaptation (UDA) is developed, which proposes to learn from both labeled source domain and unlabeled target domain [12, 40, 39].

The existing UDA works [37, 61, 33, 17, 40] usually fo-




Figure 1. (Left) Illustration of 4-class latent vectors (points) aligned in a poset formation in \mathbb{R}^2 . For a vector \mathbf{c}_k , its adjoined horizontal and vertical lines specify the feasible quadrant where its superiors \mathbf{c}_{k+i} for $i > 0$ can be positioned. (Right) Conditional spacing model $p(\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3) = p(\mathbf{c}_1)p(\mathbf{c}_2|\mathbf{c}_1)p(\mathbf{c}_3|\mathbf{c}_2)$ for 3-class. The adopted Gaussian conditional densities (4–7) and the constraints (8) enforces the sample from $p(\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_K)$ satisfy $\mathbf{c}_1 \leq \mathbf{c}_2 \leq \dots \leq \mathbf{c}_K$ with high probability [48, 10, 23].

cus on the classification or segmentation tasks without considering the inter-class correlations. However, tasks with discrete and successively labels are commonly seen in the real world. For instance, the Diabetic Retinopathy (DR) has five labels corresponding to different severity levels: 0→normal, 1→mild, 2→moderate, 3→severe, and 4→proliferative DR. In ordinal setting, class $y = 1$ is closer to $y = 2$ than $y = 4$. The misclassification of $y = 1$ to $y = 2$ or $y = 4$ may lead to different severity of misdiagnosis. The ordinal labeling system is widely adopted in medical diagnose, face age groups, and facial expression intensity, etc. Limited by the collecting cost, the UDA has great potential in the medical area, while the UDA for the ordinal label is under-presented.

A promising solution for UDA is to jointly learn the disentanglement of class-related/unrelated factors and the adaptation on the class-relevant latent space exclusively [60]. The well-performed model is expected to extract the class label informative and domain invariant features [61]. However,

the disentanglement and adaptation of ordinal factors can be challenging [2]. The *i.i.d.* Gaussian distribution sets are usually adopted as the prior of both disentangled class-related and unrelated latent space [2], which inherently postulates that the class is independently categorical.

In contrast to the *i.i.d.* categorical assumption in the conventional disentanglement framework MLVAE [2] with the vanilla Gaussian prior in VAEs [24], we propose a different setting that the extracted class-related factor \mathbf{c} is also ordinal-valued w.r.t. the corresponding ordinal label y . In other words, the value of class-related factor \mathbf{c} are also ordered.

We note that if y is an ordinal class label, i.e., $y \in \{1 < 2 < \dots < K\}$, we would expect embed this ordinal structure to form the ordinal class-related latent space. Specifically, for two instances whose class labels are close to each other, their respective extracted ordinal class-related latent vectors \mathbf{c}_1 and \mathbf{c}_2 are expected to be approximate, and vice versa [23]. For the triplet $y^1 < y^2 < y^3$, we have:

$$\|\mathbf{c}_1 - \mathbf{c}_3\| > \max\{\|\mathbf{c}_1 - \mathbf{c}_2\|, \|\mathbf{c}_2 - \mathbf{c}_3\|\}, \quad (1)$$

which aligns the class-related latent vectors and the ordinal class labels, thus yielding a model following the true data causal inference.

However, it is non-trivial to induce the ordinal inter-class correlation to the disentangled class-related latent space. Actually, directly enforcing the triplet-constraints in Eq. 1 as regularization-term for the variational lower bound of disentanglement framework is almost intractable. For the D -dimensional \mathbf{c} , there is $\mathcal{O}(D^2)$ inequalities for every pair of $(\mathbf{c}_i, \mathbf{c}_j)$, and for each inequality, we may have to introduce a slack variable to be optimized [23]. Moreover, this heuristic regularization cannot guarantee Eq. 1 is satisfied in the embedded latent space.

Motivated by the aforementioned insights, we adopted an effective and principled manner for ordinal class-related factor modeling by constructing an appropriate prior for the extracted class-related latent space. Rather than the *i.i.d.* standard Gaussian prior [24] to the non-ordinal disentanglement methods [2], our prior is constructed by a recursively conditional Gaussian (RCG) set, which explicitly imposes a restrictive partially ordered set (poset) [48, 10, 23] constraint. Since the joint prior is Gaussian, although with a full covariance, inheriting the closed form of the KL divergence term in the variational objective. Moreover, the prior can be fully factorized over the latent dimensions, which make the framework computational tractable. Furthermore, the number of parameters in the adapted prior needed to enforce the constraint for all of the pairs in a K class task can be $\mathcal{O}(D \cdot K)$, further reinforcing the model’s scalability [23]. This prior is able to assign negligible density that violates the poset constraint by a simple “three-sigma rule”. For the noisy labeled dataset, the confidence interval can be flexibly adjusted to achieve more conservative ordinal modeling.

Our contributions are summarized as follows:

- We propose to investigate the UDA of ordinal classification in a joint disentanglement and adaptation framework.
- An effective and principled ordinal prior for class-related latent space is constructed by a recursively conditional Gaussian (RCG) set [48, 10, 23]. More appealingly, we can adjust the ratio of poset violation to adapt the noise of the ordinal label. The closed form of the KL divergence with our adapted prior is computationally tractable for large-scale tasks.
- We extensively evaluate our method on both medical diagnosis (i.e., DR and CHD) and age estimation, and observe a significant improvement over the previous non-ordinal approaches in UDA.

2. Related works

Ordinal classification/regression can be processed as either multi-class classification task [14] or metric regression task [11]. The former regards the categorizes as irrelevant to each other, which ignores the inherent ordering correlation among the classes [1, 45]. The latter assumes the neighboring classes are equally distant, which can violate the non-stationary properties and easily leading to overfitting [4]. The threshold-based method proposes to define $K - 1$ binary decision boundaries [43]. Recently, the Wasserstein loss is adopted as an alternative to conventional cross-entropy loss to incorporate the inter-class correlations [30, 38, 16, 34, 31, 41, 13]. The restrictive partially ordered set (poset) has also been used to enforce the ordinal latent representation [10, 23]. The application of UDA for ordinal data is even more promising than the conventional task since the ordinal data is usually hard to label [43]. Our method is also orthogonal to the recent ordinal classification loss [30]. **Unsupervised domain adaptation** is aimed at improving the performance in an unlabeled target domain with the help of a labeled source domain [26]. The typical methods can be mainly grouped into statistic moment matching [47], domain style transfer [57, 18], feature-level adversarial learning [12, 58, 42] and self-training [61, 33, 39]. A well-performed solution is disentangling the input to class-related/unrelated factors and only aligns the class-related factor across domains [60, 3]. Moreover, the adaptation is achieved by self-training [61], which infers the pseudo-label for target samples in each iteration to gradually reduce the cross-domain variations.

However, it can be challenging to induce the ordinal prior in the disentangled latent space [30, 28, 10, 23]. We propose a joint disentangling and adaptation framework based on VAEs and tailors the prior for ordinal class-related factor (RCG) and class-unrelated factors (Gaussian).

We also tried the combination of Wasserstein loss [30] with self-training [61, 60], while the ordinal property of the to be transferred feature can not be explicitly enforced. More-

over, it is commonly challenging to combine loss-related UDA methods [50] with the ordinal loss [30].

Disentangled representation. Learning the representations that are informative to the label and robust to the other annoying variations is a primary target of representation learning [27, 36, 35, 44]. Although some works [8, 21, 7] shown that unsupervised disentanglement can be achieved by augmenting the loss of VAEs [25] also to enforce the marginal independence of the prior imposed by the encoders. While it is shown that besides sufficient inductive bias or regularity constraint is induced, the unsupervised disentanglement suffers from unidentifiability [46]. Actually, in a self-training-based UDA setting, we have the ground truth class label of source domain data, and the pseudo-label of target domain data is inferred. The group-level MLVAE model [2] is an extension of VAEs [25], which assumes a sample \mathbf{x} from class $y = k$ is a generation result of a pair of latent vectors \mathbf{c}_k and \mathbf{u} for content and instance-specific style, respectively. Therefore, all of the samples \mathbf{x} from class $y = k$ should have the same \mathbf{c}_k for content, while \mathbf{u} summarizes the other complementary variations of each sample. In contrast to MLVAE [2], we target the *ordinal-valued* \mathbf{c}_k , in which the values of content factors are *ordered* rather than categorical. Moreover, the ordinal prior induced VAEs is further incorporated in an adaptation framework.

Several recent VAE studies attempt to process the ordinal labeled data, while their objectives and setups, on the other hand, fundamentally differ from our framework. [6] assumes the ordinal paired samples are accessible for training, which is different from the ordinal classification task [43]. The variational posterior of the ordinal class is introduced in [22]. [15] utilizes the video-level label in the target domain and does not require the extracted latent vectors are able to align the ordinal constraint.

3. Methodology

In ordinal UDA, given a source domain $p(\mathbf{x}_s, y)$ and a target domain $p(\mathbf{x}_t, y)$, a labeled set \mathcal{D}_S is drawn *i.i.d* from $p(\mathbf{x}_s, y)$, and an unlabeled set \mathcal{D}_T is drawn *i.i.d* from the marginal distribution $p(\mathbf{x}_t)$. The goal of UDA is to build a good classifier in the target domain with the training on \mathcal{D}_S and \mathcal{D}_T . $y \in \{1, 2, \dots, K\}$ is the discrete and ordinal class label. In self-training UDA [61], we infer the pseudo-label \hat{y} for target samples in each iteration. Given a group of source and target images, we assume there is a ordinal prior induced shared space for the class-related factors and two domain-specific class-unrelated spaces, and the self-training [61] is applied on the shared ordinal class-related space.

For a group of source images $\{\mathbf{x}_s^n\}_{n=1}^{N_s}$ with $y = k$ and target images $\{\mathbf{x}_t^n\}_{n=1}^{N_t}$ with $\hat{y} = k$, where n indicates a group of cross-domain data instance, N_s and N_t denote the source and target domain sample numbers. Our proposed framework, is built with the VAEs [24, 2] backbones to

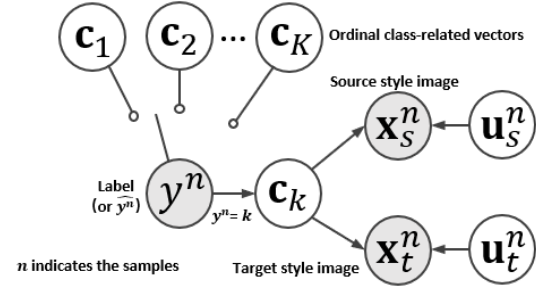


Figure 2. Graphical model representation of the proposed framework for ordinal UDA. The gray-shaded circles indicate the observed variables of the sample \mathbf{x} and its ordinal label y .

encode a sample \mathbf{x}_s^n or \mathbf{x}_t^n to the latent representation pair of $(\mathbf{c}^n, \mathbf{u}_s^n)$ or $(\mathbf{c}^n, \mathbf{u}_t^n)$, where \mathbf{c}^n is responsible for *ordinal class-related* and all of the other variations are considered as *class-unrelated feature* \mathbf{u}_s^n or \mathbf{u}_t^n for source or target domains respectively. The change of \mathbf{c}^n is expected to exclusively affect the ordinal class-related aspect and, conversely, the class-unrelated \mathbf{u}_s^n or \mathbf{u}_t^n is independent of \mathbf{c}^n . We note that the class-unrelated factors are not labeled [36].

3.1. Cross-domain Ordinal Graphical Model

By the definition of class-related \mathbf{c}^n , we expect the \mathbf{c}^n extracted from the samples within the same class should be similar or identical. Therefore, it is reasonable to calculate a representative feature \mathbf{c}_k based on a group of \mathbf{c}^n from the same class k as in [2]. Practically, we can use the average or the simple product-of-expert rule adopted in MLVAE [2].

For the well disentangled and adapted features, the observed data instance in source domain \mathbf{x}_s^n with $y^n = k$ (or target domain \mathbf{x}_t^n with $\hat{y}^n = k$) can be regarded as a generation result based on a pair of latent vectors $(\mathbf{c}_k, \mathbf{u}_s^n)$ (or $(\mathbf{c}_k, \mathbf{u}_t^n)$), respectively. We note that the ordinal label y or the pseudo label \hat{y} predicted by the self-training is always available in each iteration. This differentiates \mathbf{c}^n from $\mathbf{u}_s^n, \mathbf{u}_t^n$ in that there is one \mathbf{c}_k that governs all cross-domain instances with the same ordinal class $y = k$. However, each instance \mathbf{x}_s^n (or \mathbf{x}_t^n) have different random variable \mathbf{u}_s^n (or \mathbf{u}_t^n), which usually contains the source/target domain exclusive information and is more reasonable to configure two domain-specific VAEs.

We provide the graphical model with plate notation in Fig. 2. $p(\mathbf{x}_s^n | \mathbf{c}_k, \mathbf{u}_s^n)$ and $p(\mathbf{x}_t^n | \mathbf{c}_k, \mathbf{u}_t^n)$ are the source and target decoders, which generate images from the class-wise ordinal-related and domain-wise ordinal-unrelated latent.

3.2. Latent Priors

We configure three encoders $Enc_{s,t}^c = q_{s,t}^c(\mathbf{c}^n | \mathbf{x}_{s,t}^n)$, $Enc_s^u = q_s^u(\mathbf{u}_s^n | \mathbf{x}_s^n)$ and $Enc_t^u = q_t^u(\mathbf{u}_t^n | \mathbf{x}_t^n)$ for latent space inference, and two domain-specific decoders $Dec_s^u = p_s(\mathbf{x}_s^n | \mathbf{c}_k, \mathbf{u}_s^n)$, $Dec_t^u = p_t(\mathbf{x}_t^n | \mathbf{c}_k, \mathbf{u}_t^n)$ to generate the image with sampled latent pairs as shown in Fig. 3.




Figure 3. Our proposed ordinal UDA framework with a Recursively Conditional Gaussian (RCG) induced disentangled ordinal-related latent space. Only $Enc_{s,t}^c$ and Cls are used in testing.

For the prior of class-unrelated factors, a natural choice is the *i.i.d.* standard Gaussian $\mathcal{N}(\mathbf{0}, \mathbf{I})$. Specifically, we can set $p(\{\mathbf{u}_s^n\}_{n=1}^{N_s}) = \prod_{n=1}^{N_s} \mathcal{N}(\mathbf{u}_s^n; \mathbf{0}, \mathbf{I})$ and $p(\{\mathbf{u}_t^n\}_{n=1}^{N_t}) = \prod_{n=1}^{N_t} \mathcal{N}(\mathbf{u}_t^n; \mathbf{0}, \mathbf{I})$ following vanilla VAEs [24]. As for the class-related reference vectors, a possible choice is following the MLVAE [2] to use fully factorized standard Gaussian distributions, $p(\mathbf{c}_1, \dots, \mathbf{c}_K) = \prod_{k=1}^K \mathcal{N}(\mathbf{c}_k; \mathbf{0}, \mathbf{I})$. It should be noted that MLVAE’s *i.i.d.* prior modeling of the different class is meaningful only if the label values are categorical (i.e., discrete and independent).

Moreover, for the ordinal label, where the approximate ordinal classes are expected to be encoded. However, the *i.i.d.* Gaussian prior can be suboptimal, since it can not maintain the ordinal feature. Target for a more reasonable ordinal content prior model, we expect the triplet ordering constraint in Eq. 1 can be imposed. A possible manner of introducing this constraint is to simply adopt it as a regularization term. However, this heuristic design can not guarantee the inequity in Eq. 1, and not be scalable to a large dataset.

Following [10, 23], we propose to induce this constraint by having a *partially ordered set (poset)* [48] for the ordinal class-related latent. For each dimension $d \in \{1, \dots, D\}$, we propose to order the d_{th} element in the latent vectors, i.e., $[\mathbf{c}_1]_d < [\mathbf{c}_2]_d < \dots < [\mathbf{c}_K]_d$, where $[\mathbf{c}_k]_d$ is the d -th entry of the vector \mathbf{c}_k (i.e., $\mathbf{c}_k = [[\mathbf{c}_k]_1, [\mathbf{c}_k]_2, \dots, [\mathbf{c}_k]_d]^T$). These vectors align in a poset and satisfy the ordinal class related latent-distance constraints in Eq. 1. The case of $D = 2$ -dim latent space is shown in Fig. 1 left.

For imposing the poset constraints in the prior, we resort to the *recursively conditional Gaussian (RCG)* set [10, 23]. The framework imposes a joint Gaussian distribution over the K vectors, fully correlated for $k \in \{1, \dots, K\}$, and fully factorized for dimensions $d \in \{1, \dots, D\}$, enables the variational inferences to be tractable for computation. It is possible to effectively place negligible RCG densities that

not meet the poset constraints. Specifically, we consider a dimension-wise independent distribution,

$$p(\mathbf{c}_1, \dots, \mathbf{c}_K) = \prod_{d=1}^D p_d([\mathbf{c}_1]_d, \dots, [\mathbf{c}_K]_d), \quad (2)$$

where p_d is the density over K variables from d -th dimension of the latent vectors. We model p_d by a product of predecessor-conditioned Gaussians. For simplicity, we drop the subscript d in notation, and abuse c_k to denote $[\mathbf{c}_k]_d$, $p(c_1, \dots, c_K)$ to refer to $p_d([\mathbf{c}_1]_d, \dots, [\mathbf{c}_K]_d)$. Following [48, 10, 23], the framework can be formulated as:

$$p(c_1, \dots, c_K) = p(c_1)p(c_2|c_1) \cdots p(c_K|c_{K-1}), \quad (3)$$

where the conditionals are defined as:

$$p(c_1) = \mathcal{N}(c_1; \mu_1, \sigma_1^2) \quad (4)$$

$$p(c_2|c_1) = \mathcal{N}(c_2; \underbrace{c_1 + \Delta_2}_{:=\mu_2}, \sigma_2^2) \quad (5)$$

$$p(c_3|c_2) = \mathcal{N}(c_3; \underbrace{c_2 + \Delta_3}_{:=\mu_3}, \sigma_3^2) \quad (6)$$

\vdots

$$p(c_K|c_{K-1}) = \mathcal{N}(c_K; \underbrace{c_{K-1} + \Delta_K}_{:=\mu_K}, \sigma_K^2). \quad (7)$$

We define $p(c_k|c_{k-1})$ as a Gaussian centered at $\mu_k := c_{k-1} + \Delta_k$ with variance σ_k^2 . That is, $\Delta_k (> 0)$ is the spread between the predecessor sample c_{k-1} and the mean μ_k as shown in Fig. 1 right. We consider $\{\sigma_k, \Delta_k\}_{k=1}^K$ and μ_1 to be the free parameters of the framework that can be learned from data.

To efficiently approximate (in other words, roughly to guarantee) the poset constraint, we make each conditional distribution (pillar) separated from its adjacent neighbors through the following “three-sigma” constraint:

$$\Delta_k \geq 3\sigma_k, \quad (8)$$

which has $> 99.7\%$ probability as near certainty [53]. With Eq. 8¹, the likelihood that $c_k \leq c_{k-1}$ is negligible enforcing the desired ordering $c_1 < c_2 < \dots < c_K$. We note that the label noise is usually significant in many ordinal dataset, and the unimodal label-smoothing is widely used to achieve conservative prediction [30]. Instead of sophisticated ordinal noise modeling [28], we can simply adapt to “one/two-sigma” constraint (i.e., $\Delta_k \geq 2\sigma_k$) to tolerate the noise, which also provides more than 68% or 95% confidence interval [53].

¹The inequality constraints (8) can be easily incorporated in the conventional unconstrained optimizer modules such as PyTorch and TensorFlow through trivial reparametrizations [23], e.g., $\sigma_k := \frac{\Delta_k}{3} \text{sigmoid}(\bar{\sigma}_k)$ and $\Delta_k := \exp(\bar{\Delta}_k)$, where $\bar{\Delta}_k$ and $\bar{\sigma}_k$ are the unconstrained optimization variables, and $\text{sigmoid}(x) = 1/(1 + \exp(-x))$. Furthermore, we fix $\sigma_1 = 1$ to make the optimization numerically more stable.

The joint density for Eq. 4–7 has a closed-form. Basically, given that all of them are Gaussian, yet linear, we have the joint density $p(c_1, c_2, \dots, c_K)$ as Gaussian. Therefore, we can formulate the mean and covariance of the full joint Gaussian model as [48, 10, 23]:

$$\mathbb{E}[c_k] = \mu_1 + \Delta_2 + \dots + \Delta_k \quad (\text{for } k \geq 2) \quad (9)$$

$$\text{Cov}(c_i, c_j) = \sigma_1^2 + \dots + \sigma_{\min(i,j)}^2 \quad (10)$$

The joint distribution $p(c_1, c_2, c_3)$ For $K = 3$ is:

$$\mathcal{N}\left(\underbrace{\begin{bmatrix} \mu_1 \\ \mu_1 + \Delta_2 \\ \mu_1 + \Delta_2 + \Delta_3 \end{bmatrix}}_{:=\mathbf{a}}, \underbrace{\begin{bmatrix} \sigma_1^2 & \sigma_1^2 & \sigma_1^2 \\ \sigma_1^2 & \sigma_1^2 + \sigma_2^2 & \sigma_1^2 + \sigma_2^2 \\ \sigma_1^2 & \sigma_1^2 + \sigma_2^2 & \sigma_1^2 + \sigma_2^2 + \sigma_3^2 \end{bmatrix}}_{:=\mathbf{C}}\right), \quad (11)$$

where \mathbf{a} and \mathbf{C} are the mean vector and covariance matrix of the joint Gaussian. Plugging this back in our original prior model Eq. 2, we have:

$$p(\mathbf{c}_1, \dots, \mathbf{c}_K) = \prod_{d=1}^D \mathcal{N}([\mathbf{c}]_d; \mathbf{a}_d, \mathbf{C}_d), \quad (12)$$

where $[\mathbf{c}]_d := [[\mathbf{c}_1]_d, \dots, [\mathbf{c}_K]_d]^\top$ is the K -dim vector collecting d -th dim elements from \mathbf{c}_k . Also, \mathbf{a}_d and \mathbf{C}_d , for each $d = 1, \dots, D$, are defined by Eq. 9–11 with their own free parameters, denoted as: $(\mu_1^d, \{\Delta_k^d, \sigma_k^d\}_{k=1}^K)$ [23]. The covariance \mathbf{C}_d is not diagonal. However, the model is factorized over $d = 1, \dots, D$, the fact exploited in the next section is able to make the variational inference tractable.

3.3. Variational Inference

Given the ordinal data $\{(\mathbf{x}_s^n, y^n)\}_{n=1}^{N_s}, \{(\mathbf{x}_t^n, \hat{y}^n)\}_{n=1}^{N_t}$, we approximate the posterior by the following variational density, decomposed into the ordinal class-related and unrelated latents. The ordinal class-related posterior is further factorized over the ordinal class levels $y = 1, \dots, K$,

$$\begin{aligned} q_s(\{\mathbf{c}_k\}_{k=1}^K, \{\mathbf{u}_s^n\}_{n=1}^{N_s}) &= \prod_{k=1}^K q_{s,t}^c(\mathbf{c}_k | G_k) \prod_{n=1}^{N_s} q_s^u(\mathbf{u}_s^n | \mathbf{x}_s^n), \\ q_t(\{\mathbf{c}_k\}_{k=1}^K, \{\mathbf{u}_t^n\}_{n=1}^{N_t}) &= \prod_{k=1}^K q_{s,t}^c(\mathbf{c}_k | G_k) \prod_{n=1}^{N_t} q_t^u(\mathbf{u}_t^n | \mathbf{x}_t^n), \end{aligned} \quad (13)$$

where $q_{s,t}^c(\mathbf{c}_k | G_k)$ is short notation of $q_{s,t}^c(\mathbf{c}_k | \{\mathbf{x}_{s,t}^n\}_{n \in G_k})$ and $G_k = \{\{\mathbf{x}_s^n\}_{n=1}^{N_s}, \{\mathbf{x}_t^n\}_{n=1}^{N_t}\}$ is the set of source and target training instances with ordinal class label $y = k$ or $\hat{y} = k$. For the encoders, we adopt deep networks that take an input sample and output the means and variances of the Gaussian-distributed latents following the reparameterization trick in vanilla VAEs [24]. However, since $Enc_{s,t}^c$ requires a group of samples $\{\mathbf{x}_{s,t}^n\}_{n \in G_k}$ as its input, instead

Algorithm 1 Adaption process of our RCG induced UDA.

Initialize parameters of encoders, decoders and classifier.

repeat

Sample $G_k = \{\{\mathbf{x}_s^n\}_{n=1}^{N_s}, \{\mathbf{x}_t^n\}_{n=1}^{N_t}\}$ with $y/\hat{y} = k$.

Compute $\{\mathbf{c}^n\}_{n=1}^{N_s}, \{\mathbf{c}^n\}_{n=1}^{N_t}, \{\mathbf{u}_k^n\}_{n=1}^{N_s}, \{\mathbf{u}_k^n\}_{n=1}^{N_t}$.

Classify $\{\mathbf{c}^n\}_{n=1}^{N_s}, \{\mathbf{c}^n\}_{n=1}^{N_t}$ with *Cls*.

Calculate the loss functions: $\mathcal{L}_{s,t}^{KL}, \mathcal{L}_s^{KL}, \mathcal{L}_t^{KL}, \mathcal{L}_{s,t}^W$.

Construct \mathbf{c}_k with $\{\mathbf{c}^n\}_{n=1}^{N_s}, \{\mathbf{c}^n\}_{n=1}^{N_t}$ using Eq. 15.

Decode $\{\tilde{\mathbf{x}}_s^n\}_{n=1}^{N_s}, \{\tilde{\mathbf{x}}_t^n\}_{n=1}^{N_t}$ and calculate $\mathcal{L}_s^1, \mathcal{L}_t^1$.

//Update parameters according to gradients:

$Enc_{s,t}^c \leftarrow \mathcal{L}_{s,t}^W + \alpha \mathcal{L}_{s,t}^{KL} + \beta \mathcal{L}_s^1 + \beta \mathcal{L}_t^1 + \gamma \mathcal{L}_s^{adv} + \gamma \mathcal{L}_t^{adv}$;

$Enc_s^u \leftarrow \mathcal{L}_s^1 + \lambda \mathcal{L}_s^{KL} + \theta \mathcal{L}_s^{adv}$; $Dec_{s,t}^u \leftarrow \mathcal{L}_s^1 + \theta \mathcal{L}_s^{adv}$;

$Enc_t^u \leftarrow \mathcal{L}_t^1 + \lambda \mathcal{L}_t^{KL} + \theta \mathcal{L}_t^{adv}$; $Dec_{s,t}^u \leftarrow \mathcal{L}_t^1 + \theta \mathcal{L}_t^{adv}$;

$Cls \leftarrow \mathcal{L}_{s,t}^W$; $Dis_s \leftarrow -\mathcal{L}_s^{adv}$; $Dis_t \leftarrow -\mathcal{L}_t^{adv}$.

until deadline

of adopting a complex group encoder such as the neural statisticians [9, 20], we use a simple product-of-expert rule adopted in MLVAE [2, 23]:

$$q_{s,t}^c(\mathbf{c}_k | G_k) \propto \prod_{n \in G_k} q_{s,t}^c(\mathbf{c}^n | \mathbf{x}_{s,t}^n). \quad (14)$$

Since each $q_{s,t}^c(\mathbf{c}^n | \mathbf{x}_{s,t}^n)$ is Gaussian, the product in Eq. 14 follows a Gaussian, and it can be shown that the evidence lower bound (ELBO) can be formulated as:

$$\begin{aligned} & \sum_{k=1}^K \mathbb{E}_{q_{s,t}^c(\mathbf{c}_k | G_k)} \sum_{n \in G_k} \mathbb{E}_{q_s^u(\mathbf{u}_s^n | \mathbf{x}_s^n)} \left[\log p(\mathbf{x}_s^n | \mathbf{c}_k, \mathbf{u}_s^n) \right] \\ & + \sum_{k=1}^K \mathbb{E}_{q_{s,t}^c(\mathbf{c}_k | G_k)} \sum_{n \in G_k} \mathbb{E}_{q_t^u(\mathbf{u}_t^n | \mathbf{x}_t^n)} \left[\log p(\mathbf{x}_t^n | \mathbf{c}_k, \mathbf{u}_t^n) \right] \\ & - \sum_{n=1}^{N_s} \text{KL}\left(q_s^u(\mathbf{u}_s^n | \mathbf{x}_s^n) \parallel p(\mathbf{u}_s^n)\right) \\ & - \sum_{n=1}^{N_t} \text{KL}\left(q_t^u(\mathbf{u}_t^n | \mathbf{x}_t^n) \parallel p(\mathbf{u}_t^n)\right) \\ & - \text{KL}\left(\prod_{k=1}^K q_{s,t}^c(\mathbf{c}_k | G_k) \parallel p(\mathbf{c}_1, \dots, \mathbf{c}_K)\right). \end{aligned} \quad (15)$$

where $p(\mathbf{u}_s^n), p(\mathbf{u}_t^n) \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ as vanilla VAEs [24], while $p(\mathbf{c}_1, \dots, \mathbf{c}_K)$ is the ordinal-constrained prior given in Eq. 12. Moreover, the first two terms are the cross-domain reconstructions, which is optimized with the L1 loss $\mathcal{L}_s^1 = |\mathbf{x}_s^n - \tilde{\mathbf{x}}_s^n|, \mathcal{L}_t^1 = |\mathbf{x}_t^n - \tilde{\mathbf{x}}_t^n|$ for all samples, and the two class-unrelated KL divergences $\mathcal{L}_s^{KL}, \mathcal{L}_t^{KL}$ follow the two domain versions of MLVAE [2], with the key difference in the last term $\mathcal{L}_{s,t}^{KL}$ with induced ordinal prior RCG.




Figure 4. The distributions of average predictions on the target testset of DR UDA task. The RCG prior (orange) usually predicts more concentrated output distribution around the ground-truth label than Gaussian prior UDA (green). Best viewed in color.

Tractable ordinal latent KL-term. Despite all distributions in the KL term are Gaussians, the full dependency of \mathbf{c}_k 's over $k \in \{1, \dots, K\}$ in $p(\mathbf{c}_1, \dots, \mathbf{c}_K)$ can be ill-posed if both D and K are large [48, 10, 23]. For the $D \cdot K$ dimension input distribution, the Cholesky decomposition of $(D \cdot K \times D \cdot K)$ covariance matrix, required for computing the KL divergence, might be prohibitive if D and/or K are large. While along the latent dimensions $d = 1, \dots, D$, the prior distribution $p(\mathbf{c}_1, \dots, \mathbf{c}_K)$ is factorized as in Eq. 12. Therefore for a dimension-wise factorized encoder model, $q_{s,t}^c(\mathbf{c}_k|G_k) = \prod_{d=1}^D q_{s,t}^c([\mathbf{c}_k]_d|G_k)$, a typical setting for VAEs, we are able to reduce the complexity from $\mathcal{O}((D \cdot K)^3)$ to $\mathcal{O}(D \cdot K^3)$ [48, 10, 23]. The formal last term in Eq. 15 can be:

$$\sum_{d=1}^D \text{KL} \left(\prod_{k=1}^K q_{s,t}^c([\mathbf{c}_k]_d|G_k) \parallel p_d([\mathbf{c}_1]_d, \dots, [\mathbf{c}_K]_d) \right). \quad (16)$$

We denote $\prod_{k=1}^K q_{s,t}^c([\mathbf{c}_k]_d|G_k)$ as $\mathcal{N}(\mathbf{m}_d, \mathbf{S}_d)$, with \mathbf{S}_d diagonal by definition. Each summand in Eq. 16 is a KL divergence between Gaussians and can be written as:

$$\frac{1}{2} \left(\text{Tr}(\mathbf{C}_d^{-1} \mathbf{S}_d) + (\mathbf{a}_d - \mathbf{m}_d)^\top \mathbf{C}_d^{-1} (\mathbf{a}_d - \mathbf{m}_d) + \log \frac{|\mathbf{C}_d|}{|\mathbf{S}_d|} \right), \quad (17)$$

The computation of inverse and determinant of \mathbf{C}_d are tractable² as K is usually not a large value³.

3.4. Training processing

The pseudo-labels $\{\hat{y}^n\}_{n=1}^{N_t}$ in target domain are iteratively selected from the reliable ordinal level predictions [61]. Then the model is refined using the pseudo-labeled target images. Specifically, we sample k -class samples from both source domain ($y = k$) and target domain ($\hat{y} = k$) as a group G_k . We apply the conventional cross-entropy classification loss $\mathcal{L}_{s,t}^{CE}$ or the well-performed Wasserstein loss $\mathcal{L}_{s,t}^W$ [30] for the ordinal classifier prediction and the ground truth label or pseudo-label. It usually plays the most important role since we are focusing on the discriminative ordinal classification. We explicitly enforce the latent features to approaching a prior distribution as VAEs.

²The `inverse()` and `cholesky()` functions in PyTorch is used, which also allow auto-differentiations.

³E.g., $K = 5$ for DR labeling system.

Specifically, we minimize the KL divergence $\mathcal{L}_{s,t}^{KL}$ between $\{\mathbf{c}^n\}_{n=1}^{N_s}$, $\{\mathbf{c}^n\}_{n=1}^{N_t}$ and RCG, and \mathcal{L}_s^{KL} or \mathcal{L}_s^{KL} between $\{\mathbf{u}_k^n\}_{n=1}^{N_s}$ or $\{\mathbf{u}_k^n\}_{n=1}^{N_t}$ with Gaussian. The cross domain cycle reconstruction is utilized to enforce the disentanglement of class-related/unrelated factors, and the consistent of \mathbf{c}_k across domains. The L1 loss (\mathcal{L}_s^1 or \mathcal{L}_t^1) and adversarial loss ($\mathcal{L}_s^{adv} = \mathbb{E}[\log \text{Dis}_s(\mathbf{x}_s^n) + \log \text{Dis}_s(\tilde{\mathbf{x}}_s^n)]$ or $\mathcal{L}_t^{adv} = \mathbb{E}[\log \text{Dis}_t(\mathbf{x}_t^n) + \log \text{Dis}_t(\tilde{\mathbf{x}}_t^n)]$) are minimized to enforce the generated images to be similar as the real images in source or target domain, respectively [60, 3].

The training processing is detailed in Algorithm 1. Overall, our disentangled adaptation with self-training encourages the shared encoder $Enc_{s,t}^c$ to extract discriminative and domain-invariant representations, which can generalize and facilitate ordinal classification in the target domain.

4. Experiments

We evaluate on real-world Diabetic Retinopathy (DR) and congenital heart disease (CHD) diagnosis tasks and demonstrate its generalizability for facial age estimation tasks. We use the convolutional layers and the first fully-connected layer of the backbone network as our feature extractors, and the remaining fully-connected layers are adopted as the classifier. The decoder has the mirror structure as encoders. We implement our methods on a V100 GPU using the PyTorch and set $\alpha, \lambda, \theta = 1$, $\beta, \gamma = 0.5$ for both tasks. In our experiments, the performance is not sensitive to these hyperparameters for a relatively large range.

We use $+\mathcal{L}_{s,t}^W$ to denote the solutions using the Wasserstein loss [30] rather than the cross-entropy loss, while $-\mathcal{L}_{s,t}^{adv}$ indicates the ablation of the adversarial loss. 2σ or 3σ indicates the "two-sigma rule" or "three-sigma rule" version of RCG, respectively. For the supervised setting, we use pre-train the model with the labeled source domain, and fine-tune the model with labeled target samples as the LzyUCNN [52].

4.1. Diabetic Retinopathy Diagnosis

Two public available ordinal DR datasets are adopted for DNN implementations. The Kaggle Diabetic Retinopathy (KDR)⁴ is utilized as our source domain, and the recent In-

⁴<https://www.kaggle.com/c/diabetic-retinopathy-detection>

Method	Accuracy (%) \uparrow	QWK \uparrow	MAE \downarrow
Source only	48.6 \pm 0.03	45.8 \pm 0.02	0.61 \pm 0.01
CRST [61]	56.5 \pm 0.01	53.2 \pm 0.01	0.50 \pm 0.02
TPN [50]	56.8 \pm 0.02	52.7 \pm 0.01	0.52 \pm 0.01
DMRL [59]	57.3 \pm 0.01	53.0 \pm 0.03	0.50 \pm 0.02
DG-Net++ [60]	57.0 \pm 0.02	53.6 \pm 0.02	0.51 \pm 0.03
CRST [61] $+\mathcal{L}_{s,t}^W$ [30]	58.8 \pm 0.03	54.2 \pm 0.01	0.47 \pm 0.02
DG-Net++ [60] $+\mathcal{L}_{s,t}^W$	59.3 \pm 0.02	54.5 \pm 0.02	0.48 \pm 0.02
RCGUDA: 2σ	61.8 \pm 0.01	56.4 \pm 0.01	0.41 \pm 0.02
RCGUDA: 3σ	61.6 \pm 0.03	56.2 \pm 0.02	0.43 \pm 0.01
RCGUDA: $2\sigma-\mathcal{L}_{s,t}^{adv}$	61.5 \pm 0.02	56.1 \pm 0.02	0.42 \pm 0.02
RCGUDA: $2\sigma+\mathcal{L}_{s,t}^W$	62.3 \pm 0.02	56.9 \pm 0.01	0.40 \pm 0.01
Supervised [52]	63.1 \pm 0.02	58.7 \pm 0.02	0.38 \pm 0.01

Table 1. Experimental results for DR UDA task. \uparrow larger is better.

dian Diabetic Retinopathy Image Dataset (IDRiD) dataset [51] is used as our target domain. In both datasets, the diabetic retinal images are grouped following the International Clinical Diabetic Retinopathy Scale, where level 0 to 4 representing the No DR, Mild DR, Moderate DR, Severe DR, and Proliferative DR, respectively. The macular edema severity are labeled according to the occurrences of hard exudates near to macula center region.

The KDR consists of 88,702 left or right fundus (i.e., interior surface at the back of the eye) images from 17,563 patients with the size varies from 433×289 to 3888×2592 . Although its large scale, there is no standard field-of-view (FOV) and camera type. The images in IDRiD have the unified size (4288×2848), FOV (50), and use the same Kowa VX-10 α camera. More importantly, the IDRiD is the first dataset of Indians which has a significant shift between KDR w.r.t. population.

In the standard setting of IDRiD⁵, there are 413 images for training and 103 images for testing. We set its 413 training images as our *unlabeled* target domain, while the 103 testing images are only used for UDA testing. We preprocess and resized the image to 896×896 as [52]. The ResNet [19] model in LzyUNCC [52] has been adopted as our backbone. Of note, LzyUNCC [52] uses both *labeled* KDR and *labeled* IDRiD for *supervised* training, which can be regard as the “upper bound”.

The results of the DR task are shown in Tab. 1. Other than the accuracy and Mean Absolute Error (MAE), we also report the quadratic weighted kappa (QWK)⁶. QWK can reflect the misclassification proportional to the distance between the prediction and ground-truth [30, 28]. Several recent UDA methods are compared, e.g., self-training based CRST [61], Prototypical Network (TPN) [50], disentangle-based adversarial and self-training method (DG-Net++) [60], dual

⁵<https://iee-dataport.org/open-access/indian-diabetic-retinopathy-image-dataset-idrid>

⁶<https://www.kaggle.com/c/diabetic-retinopathy-detection/overview/evaluation>

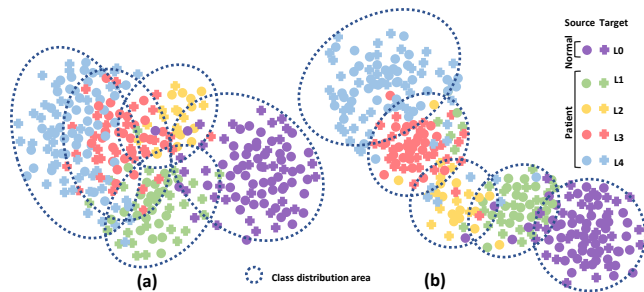


Figure 5. t-SNE visualization of the features on CHD UDA task extracted by $Enc_{s,t}^c$. Best viewed in color.

mixup regularized learning (DMRL) [59]. Although these methods have demonstrated their effectiveness in categorical classification based UDA, they do not take the inter-class correlation into account. A possible solution is applying [61, 60] on top of the Wasserstein loss [30]. We note that [50] itself is the improvement of the loss function, which is not directly compatible with the Wasserstein loss [30].

Our proposed RCGUDA outperforms these solutions consistently. Our improvements w.r.t. QWK and MAE are even more appealing. We note that QWK and MAE can be the more reasonable measure for ordinal data [43], since it considers the severity of different misclassification. In Fig. 4, we illustrated that RCGUDA could provide a more concentrated averaged prediction distribution for the target testing samples. The 2σ version can be slightly better than 3σ version, considering the label noise. We also provide the ablation study of adversarial loss, which not only improves the performance but also speeds up the convergence of about 1.2 times. Furthermore, our framework is also orthogonal to the advanced ordinal classification loss, e.g., [30]. With the help of Wasserstein loss [30], the UDA performance can approach the previously supervised LzyNUCC [52].

4.2. Real-world CHD Transfer

Congenital heart disease (CHD) is the most prevalent congenital disability and the primary cause of death in newborns. The clinical diagnosis is usually based on the selected keyframes from the five-view echocardiogram videos. We collect a multi-center dataset of five-view echocardiograms of 1,600 labeled source subjects (L0-L4 levels) from Beijing Children’s Hospital and 800 unlabeled target subjects from the Harvard Medical School. To quantify the severity levels, each subject is confirmed by either at least two senior ultrasound doctors or final intraoperative diagnosis. Note that there are discrepancies w.r.t. imaging devices (PHILIPS iE 33 vs. EPIQ 7C), patient populations, and the echocardiogram imaging experience of clinical doctors.

We adopt the standard multi-channel network for five-view echocardiogram processing. We use 80% and 20% target subjects for training and testing, respectively. For com-

Method	Accuracy (%) \uparrow	QWK \uparrow	MAE \downarrow
Source only	66.2 \pm 0.02	60.3 \pm 0.01	0.42 \pm 0.02
CRST [61]+ $\mathcal{L}_{s,t}^W$ [30]	72.5 \pm 0.01	69.9 \pm 0.02	0.36 \pm 0.02
DG-Net++ [60]+ $\mathcal{L}_{s,t}^W$	72.4 \pm 0.01	70.2 \pm 0.01	0.35 \pm 0.02
RCGUDA:2 σ	74.6 \pm 0.02	71.8 \pm 0.01	0.33 \pm 0.01
RCGUDA:3 σ	75.0 \pm 0.03	72.1 \pm 0.02	0.32 \pm 0.03
RCGUDA:3 σ - $\mathcal{L}_{s,t}^{adv}$	74.5 \pm 0.01	71.9 \pm 0.03	0.33 \pm 0.01
RCGUDA:3 σ + $\mathcal{L}_{s,t}^W$	75.6 \pm 0.03	72.5 \pm 0.02	0.31 \pm 0.01
Supervised+ $\mathcal{L}_{s,t}^W$ [30]	78.4 \pm 0.02	74.6 \pm 0.01	0.28 \pm 0.02

Table 2. Experimental results for CHD UDA task. \uparrow larger is better.

parison, we re-implement the recent state-of-the-art methods with the same backbone and experiment setting. The results are shown in Tab. 2. The improvements over the previous state-of-the-art are consistent with the DR UDA task. The 3 σ performs better than 2 σ , since the relatively accurate label.

In Fig. 5 left, we can see that the CHD samples tend to be distributed unordered. With our RCG prior, both the samples are ordinally grouped into the high-density region w.r.t. severity level in Fig. 5 right. The ordinal prior can be effectively induced into the latent space features.

4.3. Face age estimation

Although our framework has great potentials for medical image analysis, it can be generalized to the facial age estimation, which also has an ordinal label. We use MORPH Album II [54] as our labeled source domain, and choose IMDB-WIKI [55, 56] as unlabeled source domain.

MORPH Album II [54] is a widely used dataset for estimating face age. It collects 55,134 RGB images from 13,617 people. The age is labeled between 16 to 77. Besides White or Black, the proportion of the other ethnicity is relatively low (4%). IMDB-WIKI dataset⁷ is consisted of the celebrity faces from the IMDB and Wikipedia. Each image is labeled as the difference between the published data and the birthday of the account. In summary, the age is ranged from 5 to 90 years old. There are both gray and RGB images, and stamp portraits, which lead to the age label can be quite noisy and biased. Following the evaluation protocol, the face detector is applied to crop the face region for subsequent processing, and the multi-face images are dropped. Then, we have 224,418 images out of the original 523,051 images. We divide the IMDB-WIKI dataset following the subject ID-independent manner and use 50% for the unlabeled target domain training data and 50% for testing.

We configure the age as the expected ordinal content factor. To unify the age groups in two datasets, we partition age values into 6 age-range classes. A age-range class involves 10 years interval, i.e., 16-25, 26-35, 36-45, 46-55, 56-65, 66-75. We note that the variations other than age, e.g., pose, illumination and facial expression, are regarded as ordinal-

⁷<https://data.vision.ee.ethz.ch/cvl/rrothe/imdb-wiki/>

Method	Accuracy (%) \uparrow	QWK \uparrow	MAE \downarrow
Source only	48.7 \pm 0.01	42.8 \pm 0.02	0.46 \pm 0.01
CRST [61]+ $\mathcal{L}_{s,t}^W$ [30]	61.3 \pm 0.02	60.1 \pm 0.02	0.39 \pm 0.02
DG-Net++ [60]+ $\mathcal{L}_{s,t}^W$	62.4 \pm 0.02	60.7 \pm 0.01	0.39 \pm 0.01
RCGUDA:3 σ	67.8 \pm 0.02	63.2 \pm 0.02	0.36 \pm 0.01
RCGUDA:2 σ	68.1 \pm 0.02	64.5 \pm 0.03	0.35 \pm 0.02
RCGUDA:2 σ - $\mathcal{L}_{s,t}^{adv}$	67.7 \pm 0.03	63.0 \pm 0.03	0.36 \pm 0.01
RCGUDA:2 σ + $\mathcal{L}_{s,t}^W$	68.3 \pm 0.02	64.8 \pm 0.01	0.34 \pm 0.02
Supervised+ $\mathcal{L}_{s,t}^W$ [30]	72.6 \pm 0.03	70.4 \pm 0.02	0.32 \pm 0.01

Table 3. Experimental results for age UDA task. \uparrow larger is better.

class unrelated (unlabeled). Following [49, 28], we choose the VGG16 as our backbone. The batch size was set to 64 for VGG16. In addition, the initial learning rate is 0.001. We reduce the learning rate by multiplying 0.1 per 15 epochs.

The results are shown in Tab. 3. The improvement of our RCGUDA is consistent with the medical tasks. We note that the age group in IMDB-WIKI is significantly noisier than the medical data, and demonstrates large inner-class variations. Therefore, the "two-sigma rule" version of RCG can perform better than its "three-sigma rule" version. Actually, we can also choose the decimal factor to sigma, while the performance can be stable for a large range and have the same QWK from 1.8 σ to 2.4 σ version of RCG.

5. Conclusions

This paper targets to explicitly induce the ordinal prior for ordinal classification UDA. We propose to meet the ordinal constraint by defining a partially ordered set (poset). Specifically, we developed a recursively conditional Gaussian (RCG) for ordered constraint modeling, which admits a tractable joint distribution prior. By enforcing RCG as the prior distribution of our joint disentanglement and adaptation framework rather than the conventional Gaussian, the performance on DR and CHD diagnosis and facial age estimation tasks can be significantly improved. The three/two-sigma setting of RCG can flexibly fit for the different degrees of label noise without the sophisticated ordinal noisy modeling. It is also orthogonal to the recently developed ordinal classification loss, which can be simply added to our RCGUDA. Our experiments evidenced that our RCG can be a powerful and versatile unsupervised solution for ordinal UDA. For future work, we plan to explore the task with a more sophisticated ordinal label, e.g., pathology level segmentation.

Acknowledgments

The funding support from National Institute of Health (NIH), USA (NS061841, NS095986), Fanhan Tech., Jiangsu NSF (BK20200238), PolyU Central Research Grant (G-YBJW) and Hong Kong Government General Research Fund GRF (152202/14E) are greatly appreciated.

References

- [1] Christopher Beckham and Christopher Pal. Unimodal probability distributions for deep ordinal classification. *arXiv preprint arXiv:1705.05278*, 2017. [2](#)
- [2] Diane Bouchacourt, Ryota Tomioka, and Sebastian Nowozin. Multi-level variational autoencoder: Learning disentangled representations from grouped observations. In *The Thirty-Second AAAI Conference on Artificial Intelligence*, 2018. [2](#), [3](#), [4](#), [5](#)
- [3] Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain separation networks. *arXiv preprint arXiv:1608.06019*, 2016. [2](#), [6](#)
- [4] Kuang-Yu Chang, Chu-Song Chen, and Yi-Ping Hung. Ordinal hyperplanes ranker with cost sensitivities for age estimation. In *Computer vision and pattern recognition (cvpr), 2011 IEEE conference on*, pages 585–592. IEEE, 2011. [2](#)
- [5] Tong Che, Xiaofeng Liu, Site Li, Yubin Ge, Ruixiang Zhang, Caiming Xiong, and Yoshua Bengio. Deep verifier networks: Verification of deep discriminative models with deep generative models. *AAAI*. [1](#)
- [6] Junxiang Chen and Kayhan Batmanghelich. Robust ordinal VAE: Employing noisy pairwise comparisons for disentanglement, 2019. [arXiv:1910.05898](#). [3](#)
- [7] Ricky T. Q. Chen, Xuechen Li, Roger Grosse, and David Duvenaud. Isolating sources of disentanglement in variational autoencoders. In *Advances in Neural Information Processing Systems*, 2018. [3](#)
- [8] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel. InfoGAN: Interpretable representation learning by information maximizing Generative Adversarial Nets, 2016. In *Advances in Neural Information Processing Systems*. [3](#)
- [9] Harrison Edwards and Amos Storkey. Towards a neural statistician, 2017. In *Proceedings of the Second International Conference on Learning Representations (ICLR)*. [5](#)
- [10] Marco Fattore. Partially ordered sets and the measurement of multidimensional ordinal deprivation. *Social Indicators Research*, 128(2):835–858, 2016. [1](#), [2](#), [4](#), [5](#), [6](#)
- [11] Yun Fu and Thomas S Huang. Human age estimation with regression on discriminative aging manifold. *IEEE Transactions on Multimedia*, 10(4):578–584, 2008. [2](#)
- [12] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 2016. [1](#), [2](#)
- [13] Yubin Ge, Site Li, Xuyang Li, Fangfang Fan, Wanqing Xie, Jane You, and Xiaofeng Liu. Embedding semantic hierarchy in discrete optimal transport for risk minimization. In *ICASSP*, pages 2835–2839. IEEE, 2021. [2](#)
- [14] Xin Geng, Zhi-Hua Zhou, and Kate Smith-Miles. Automatic age estimation based on facial aging patterns. *IEEE Transactions on pattern analysis and machine intelligence*, 29(12):2234–2240, 2007. [2](#)
- [15] Eric Granger, Patrick Cardinal, et al. Deep domain adaptation for ordinal regression of pain intensity estimation using weakly-labelled videos. *arXiv preprint arXiv:2008.06392*, 2020. [3](#)
- [16] Yuzhuo Han, Xiaofeng Liu, Zhenfei Sheng, Yutao Ren, Xu Han, Jane You, Risheng Liu, and Zhongxuan Luo. Wasserstein loss-based deep object detection. In *CVPR Workshops*, pages 998–999, 2020. [2](#)
- [17] Gewen He, Xiaofeng Liu, Fangfang Fan, and Jane You. Classification-aware semi-supervised domain adaptation. In *CVPR Workshops*, pages 964–965, 2020. [1](#)
- [18] Gewen He, Xiaofeng Liu, Fangfang Fan, and Jane You. Image2audio: Facilitating semi-supervised audio emotion recognition with facial expression image. In *CVPR Workshops*, 2020. [2](#)
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [7](#)
- [20] Luke B. Hewitt, Maxwell I. Nye, Andreea Gane, Tommi Jaakkola, and Joshua B. Tenenbaum. The variational homocoder: Learning to learn high capacity generative models from few examples, 2018. *Uncertainty in Artificial Intelligence (UAI)*. [5](#)
- [21] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner. Beta-VAE: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017. [3](#)
- [22] Joel Jaskari and Jyri J. Kivinen. A novel variational autoencoder with applications to generative modelling, classification, and ordinal regression, 2018. [arXiv:1812.07352](#). [3](#)
- [23] Minyoung Kim and Vladimir Pavlovic. Ordinal-content vae: Isolating ordinal-valued content factors in deep latent variable models. In *2020 IEEE International Conference on Data Mining (ICDM)*, pages 252–261. IEEE, 2020. [1](#), [2](#), [4](#), [5](#), [6](#)
- [24] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. [2](#), [3](#), [4](#), [5](#)
- [25] Diederik P. Kingma and Max Welling. Auto-encoding variational Bayes, 2014. In *Proceedings of the Second International Conference on Learning Representations, ICLR*. [3](#)
- [26] Wouter M Kouw. An introduction to domain adaptation and transfer learning. *arXiv*, 2018. [2](#)
- [27] Xiaofeng Liu, Yang Chao, Jane J You, C-C Jay Kuo, and Bhagavatula Vijayakumar. Mutual information regularized feature-level frankenstein for discriminative recognition. *IEEE T-PAMI*, 2021. [3](#)
- [28] Xiaofeng Liu, Fangfang Fan, Lingsheng Kong, Zihui Diao, Wanqing Xie, Jun Lu, and Jane You. Unimodal regularized neuron stick-breaking for ordinal classification. *Neurocomputing*, 2020. [2](#), [4](#), [7](#), [8](#)
- [29] Xiaofeng Liu, Zhenhua Guo, Site Li, Fangxu Xing, Jane You, C-C Jay Kuo, Georges El Fakhri, and Jonghye Woo. Adversarial unsupervised domain adaptation with conditional and label shift: Infer, align and iterate. *ICCV*, 2021. [1](#)
- [30] Xiaofeng Liu, Xu Han, Yukai Qiao, Yi Ge, Site Li, and Jun Lu. Unimodal-uniform constrained wasserstein training for medical diagnosis. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019. [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)

- [31] Xiaofeng Liu, Yuzhuo Han, Song Bai, Yi Ge, Tianxing Wang, Xu Han, Site Li, Jane You, and Jun Lu. Importance-aware semantic segmentation in self-driving with discrete wasserstein training. In *AAAI*, volume 34, pages 11629–11636, 2020. **2**
- [32] Xiaofeng Liu, Bo Hu, Linghao Jin, Xu Han, Fangxu Xing, Jinsong Ouyang, Jun Lu, Georges El Fakhri, and Jonghye Woo. Domain generalization under conditional and label shifts via variational bayesian inference. In *IJCAI*, 2021. **1**
- [33] Xiaofeng Liu, Bo Hu, Xiongchang Liu, Jun Lu, Jane You, and Lingsheng Kong. Energy-constrained self-training for unsupervised domain adaptation. In *ICPR*, 2021. **1, 2**
- [34] Xiaofeng Liu, Wenxuan Ji, Jane You, Georges El Fakhri, and Jonghye Woo. Severity-aware semantic segmentation with reinforced wasserstein training. In *CVPR*, pages 12566–12575, 2020. **2**
- [35] Xiaofeng Liu, Linghao Jin, Xu Han, and Jane You. Mutual information regularized identity-aware facial expression recognition in compressed video. *Pattern Recognition*, 2021. **3**
- [36] Xiaofeng Liu, Site Li, Lingsheng Kong, Wanqing Xie, Ping Jia, Jane You, and BVK Kumar. Feature-level frankenstein: Eliminating variations for discriminative recognition. In *CVPR*, 2019. **3**
- [37] Xiaofeng Liu, Xiongchang Liu, Bo Hu, Wenxuan Ji, Fangxu Xing, Jun Lu, Jane You, C-C Jay Kuo, Georges El Fakhri, and Jonghye Woo. Subtype-aware unsupervised domain adaptation for medical diagnosis. In *AAAI*, volume 35, pages 2189–2197, 2021. **1**
- [38] Xiaofeng Liu, Yunhong Lu, Xiongchang Liu, Song Bai, Site Li, and Jane You. Wasserstein loss with alternative reinforcement learning for severity-aware semantic segmentation. *IEEE Transactions on Intelligent Transportation Systems*, 2020. **2**
- [39] Xiaofeng Liu, Fangxu Xing, Maureen Stone, Jiachen Zhuo, Reese Timothy, Jerry L Prince, Georges El Fakhri, and Jonghye Woo. Generative self-training for cross-domain unsupervised tagged-to-cine mri synthesis. *MICCAI*, 2021. **1, 2**
- [40] Xiaofeng Liu, Fangxu Xing, Chao Yang, Georges El Fakhri, and Jonghye Woo. Adapting off-the-shelf source segmenter for target medical image segmentation. *MICCAI*, 2021. **1**
- [41] Xiaofeng Liu, Yang Zou, Tong Che, Peng Ding, Ping Jia, Jane You, and BVK Kumar. Conservative wasserstein training for pose estimation. In *ICCV*, pages 8262–8272, 2019. **2**
- [42] Xiaofeng Liu, Yang Zou, Lingsheng Kong, Zhihui Diao, Junliang Yan, Jun Wang, Site Li, Ping Jia, and Jane You. Data augmentation via latent space interpolation for image classification. In *ICPR*. IEEE, 2018. **2**
- [43] Xiaofeng Liu, Yang Zou, Yuhang Song, Chao Yang, Jane You, and BVK Vijaya Kumar. Ordinal regression with neuron stick-breaking for medical diagnosis. In *European Conference on Computer Vision*, pages 335–344. Springer, 2018. **2, 3, 7**
- [44] Yanzhu Liu, Adams Wai-Kin Kong, and Chi Keong Goh. Deep ordinal regression based on data relationship for small datasets. In *IJCAI*, pages 2372–2378, 2017. **3**
- [45] Yanzhu Liu, Adams Wai Kin Kong, and Chi Keong Goh. A constrained deep neural network for ordinal regression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 831–839, 2018. **2**
- [46] Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Rätsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging Common Assumptions in the Unsupervised Learning of Disentangled Representations, 2019. arXiv:1811.12359. **3**
- [47] Long, Cao, Wang, and Jordan. Conditional adversarial domain adaptation. In *NIPS*, 2018. **2**
- [48] Holbrook Mann MacNeille. Partially ordered sets. *Transactions of the American Mathematical Society*, 42(3):416–460, 1937. **1, 2, 4, 5, 6**
- [49] Hongyu Pan, Hu Han, Shiguang Shan, and Xilin Chen. Mean-variance loss for deep age estimation from a face. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5285–5294, 2018. **8**
- [50] Yingwei Pan, Ting Yao, Yehao Li, Yu Wang, Chong-Wah Ngo, and Tao Mei. Transferrable prototypical networks for unsupervised domain adaptation. In *CVPR*, 2019. **3, 7**
- [51] Prasanna Porwal, Samiksha Pachade, Ravi Kamble, Manesh Kokare, Girish Deshmukh, Vivek Sahasrabudhe, and Fabrice Meriaudeau. Indian diabetic retinopathy image dataset (idrid): a database for diabetic retinopathy screening research. *Data*, 3(3):25, 2018. **7**
- [52] Prasanna Porwal, Samiksha Pachade, Manesh Kokare, Girish Deshmukh, Jaemin Son, Woong Bae, Lihong Liu, Jianzong Wang, Xinhui Liu, Liangxin Gao, et al. Idrid: Diabetic retinopathy–segmentation and grading challenge. *Medical image analysis*, 59:101561, 2020. **6, 7**
- [53] Friedrich Pukelsheim. The three sigma rule. *The American Statistician*, 48(2):88–91, 1994. **4**
- [54] Karl Ricanek and Tamirat Tesafaye. Morph: A longitudinal image database of normal adult age-progression. In *Automatic Face and Gesture Recognition, 2006. FGR 2006. 7th International Conference on*, pages 341–345. IEEE, 2006. **8**
- [55] Rasmus Rothe, Radu Timofte, and Luc Van Gool. Dex: Deep expectation of apparent age from a single image. In *IEEE International Conference on Computer Vision Workshops (IC-CVW)*, December 2015. **8**
- [56] Rasmus Rothe, Radu Timofte, and Luc Van Gool. Deep expectation of real and apparent age from a single image without facial landmarks. *IJCV*, July 2016. **8**
- [57] Swami Sankaranarayanan, Yogesh Balaji, Carlos D Castillo, and Rama Chellappa. Generate to adapt: Aligning domains using generative adversarial networks. In *CVPR*, 2018. **2**
- [58] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, 2017. **2**
- [59] Yuan Wu, Diana Inkpen, and Ahmed El-Roby. Dual mixup regularized learning for adversarial domain adaptation. *ECCV*, 2020. **7**
- [60] Yang Zou, Xiaodong Yang, Zhiding Yu, BVK Kumar, and Jan Kautz. Joint disentangling and adaptation for cross-domain person re-identification. *ECCV*, 2020. **1, 2, 6, 7, 8**
- [61] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *ICCV*, 2019. **1, 2, 3, 6, 7, 8**